
Feature-Critic Networks for Heterogeneous Domain Generalisation

Yiyi Li^{*12} Yongxin Yang^{*2} Wei Zhou¹ Timothy M. Hospedales²³

Abstract

The well known domain shift issue causes model performance to degrade when deployed to a new target domain with different statistics to training. Domain adaptation techniques alleviate this, but need some instances from the target domain to drive adaptation. Domain generalisation is the recently topical problem of learning a model that generalises to unseen domains out of the box, and various approaches aim to train a domain-invariant feature extractor, typically by adding some manually designed losses. In this work, we propose a *learning to learn* approach, where the auxiliary loss that helps generalisation is itself learned. Beyond conventional domain generalisation, we consider a more challenging setting of *heterogeneous* domain generalisation, where the unseen domains do not share label space with the seen ones, and the goal is to train a feature representation that is useful off-the-shelf for novel data and novel categories. Experimental evaluation demonstrates that our method outperforms state-of-the-art solutions in both settings.

1. Introduction

A shift in data statistics between training and testing is often unavoidable in real-world applications, and leads to a significant negative impact on the performance of machine learning models in practice. This motivates research into methods to ameliorate the impact of domain shift, including Domain Adaption (DA) (Bousmalis et al., 2016; Ganin & Lempitsky, 2015; Long et al., 2015; 2016) and Domain Generalisation (DG) (Muandet et al., 2013; Ghifary et al., 2015; Li et al., 2018a; Shankar et al., 2018).

^{*}Equal contribution ¹College of Computer, National University of Defense Technology, Hunan, China ²School of Informatics, The University of Edinburgh, Edinburgh, UK ³Samsung AI Centre, Cambridge, UK. Correspondence to: Timothy M. Hospedales <t.hospedales@ed.ac.uk>.

Unsupervised Domain Adaptation (UDA) (Long et al., 2016; Saito et al., 2017; Shu et al., 2018) methods operate in the setting where we can access unlabelled testing (target) domain data during training to drive model adaptation and compensate for the domain shift. Domain Generalisation addresses the harder setting, where a model trained on a set of source domains should perform well on a novel target domain with different data statistics, without requiring any access to target domain data during training. That is, the model should be robust enough out-of-the-box to perform well in a new domain, without further parameter updates. Both DA and DG methods almost always assume the label space is consistent across both source and target domains.

In the case of disjoint label spaces between source and target domain, we term the domain generalisation problem as one of *heterogeneous* domain generalisation. In this case a *feature representation* trained on a source domain should generalise to supporting recognition of novel categories in a novel target domain. This problem setting is actually widely encountered. The central example is the ubiquitous computer vision pipeline where a CNN feature extractor pre-trained on ImageNet is re-used for diverse applications. If data, computation, and human expert time is available, the feature can be fine-tuned on the target problem. However, for many practical applications lacking one or more of these requirements, standard practice is to use an ImageNet CNN off-the-shelf as a fixed feature extractor, and train a shallow model such as SVM or KNN for the new problem (Donahue et al., 2014; Razavian et al., 2014). This pipeline is an example of the heterogeneous domain generalisation setting, in that a feature is being asked to generalise to supporting recognition of novel categories in data with novel statistics. The ImageNet pre-trained feature is strong enough to do a reasonable job of this already. However, given the ubiquity of this pipeline, providing an improved general purpose feature would be widely beneficial. In this paper we aim to do exactly this by presenting a novel method that explicitly trains a feature to prepare it for domain and label shift. We demonstrate this via performing heterogeneous DG on the Visual Decathlon benchmark (Rebuffi et al., 2017). This also provides the largest scale evaluation of DG to date.

We are inspired by recent meta-learning learning methods that perform episodic training (Finn et al., 2017; Snell et al., 2017; Ravi & Larochelle, 2017) to simulate the train/test

process to improve few-shot learning. In this work, we propose to perform meta-learning to improve feature extractor training, and deliver a better model for both homogeneous and heterogeneous DG problems.

To realise our idea, we simulate training-to-testing domain shift by splitting our source domains into virtual training and testing (i.e., validation) domains. The source model is decomposed into feature extractor and task networks (i.e., a classifier network in our case). Crucially we then introduce a feature-critic network that learns to criticise the quality of the features produced by the feature network, specifically with regards to their robustness to the simulated domain shift. This feature-critic provides a learned auxiliary loss which provides an additional source of feedback to the feature network (besides the conventional supervised classification loss via the task network), and enables it to produce a more robust feature. The feature, task and critic networks are trained together end-to-end in a meta-learning pipeline. Our evaluation shows good performance in the conventional DG setting using Rotated MNIST (Ghifary et al., 2015; Motiian et al., 2017) and PACS (Li et al., 2017a) benchmarks, as well as the heterogeneous DG setting using the larger scale Visual Decathlon (VD) (Rebuffi et al., 2017) benchmark.

2. Related Work

Multi-Domain Learning (MDL) MDL addresses training a single model capable of solving multiple datasets (domains). If the data is relatively small and the domains are similar, this sharing can lead to improved performance compared to training a separate model per domain (Yang & Hospedales, 2015). On the other hand, for diverse domains with large data, MDL may under-perform a single model per domain; but is nonetheless of interest due to the simplicity of a single model and its better memory scalability compared to a separate model per domain (Rebuffi et al., 2017; 2018). We mention MDL here, because DG methods typically train on multiple source domains as per MDL – but furthermore aim to generalise to novel held out domains.

Domain Generalisation (DG) DG relates to domain-adaptation in that we care about performance on a target domain, rather than source domains; however it considers the case where target domain samples are unavailable during training, so the model must generalise directly rather than adapt to the target domain. DG is related to conventional generalisation: where models learned on a set of training instances generalise to novel testing instances, for example by regularisation. However it operates at a higher level, where we aim to help models trained on a set of training domains generalise to a novel testing domain.

Most existing DG approaches can be split into three categories: feature-based methods, classifier-based methods,

and data augmentation methods. Feature-based methods: These aim to generate a domain-invariant representation. For example where the distance between the empirical distributions of the source and target examples is minimized (Li et al., 2018b; Muandet et al., 2013; Li et al., 2018a). Classifier-based methods: Some aim to enhance generalisation by fusing multiple sub-classifiers learned from source domains (Duan et al., 2012; Niu et al., 2015a;b), and others learn an improved classifier regulariser using source samples – notably the recently proposed MetaReg (Balaji et al., 2018). Data augmentation methods: CrossGrad (Shankar et al., 2018) generates provides domain-guided perturbations of input instances, which are then used to train a more robust model. Volpi et al. (2018) defines an adaptive data augmentation scheme by appending adversarial examples at each iteration. Our Feature-Critic approach falls into the feature-based category, but meta-learns a feature-critic network to train a robust shared feature extractor.

Few studies have considered the heterogeneous DG setting, where the domains do not share the same label space. We do not expect the *classifier* to generalise directly to the target domain (impossible due to the change in label space), but we do aim to improve the robustness of a source-domain trained *feature* in terms of its generalisation to successfully represent a novel problem. Most existing DG methods cannot be applied here. We show how to modify MetaReg (Balaji et al., 2018) and Reptile (Nichol et al., 2018) to address this DG setting. The most relevant benchmark is Visual Decathlon (VD) (Rebuffi et al., 2017). The VD benchmark was proposed to evaluate multi-domain and lifelong (Rosenfeld & Tsotsos, 2018) learning. We re-purpose it for DG evaluation. In this case a model trained on the six largest datasets in VD should produce a feature which provides a general and robust enough encoding to allow the four smaller datasets to be classified with a simple shallow classifier.

Meta-Learning Meta-learning (a.k.a. learning to learn, (Schmidhuber et al., 1997; Thrun & Pratt, 1998)) has received resurgence in interest recently with applications in few-shot learning (Li et al., 2017b; Snell et al., 2017; Sung et al., 2018) and beyond (Xu et al., 2018). In few-shot meta-learning, a common strategy is to simulate the few-shot learning scenario by randomly drawing few-shot train/test episodes from the full training set. We adapt this episodic training strategy by creating virtual training and testing splits of our source domains in each mini-batch.

A few methods have applied related episodic meta-learning strategies in DG (Li et al., 2018a; Balaji et al., 2018). MLDG (Li et al., 2018a) defined a heuristic gradient descent update rule based on the gradients of the simulated training and testing domains. MetaReg (Balaji et al., 2018) trains the weights of the *classifier*'s regulariser so as to produce a more general classifier for a fixed feature extractor. In contrast,




Figure 1. Illustration of the Feature-Critic learning framework.

our Feature-Critic produces a more general *feature extractor* that can be used with any classifier. This is achieved by simultaneously learning an auxiliary loss function (Gygli et al., 2017; Sung et al., 2017) (i.e., the critic network) that trains the feature extractor for improved domain invariance.

3. Methodology

We introduce the proposed method under the heterogeneous DG setting, but it is straightforwardly applicable to conventional (homogeneous) DG as a special case. Assuming that we have N domains (datasets) $\mathcal{D} = \{D_1, D_2, \dots, D_N\}$, and each domain contains a set of data-label pairs, i.e., $D_i = \{X^{(i)}, Y^{(i)}\}$. We also have the training split of target (testing) domain, $D_{N+1} = \{X^{(N+1)}, Y^{(N+1)}\}$, but we can not access this for feature learning.

We assume a CNN model split into two parts: feature extractor f_θ and classifier g_ϕ . For heterogeneous DG, we have N classifiers, denoted $g_{\phi_1}, g_{\phi_2}, \dots, g_{\phi_N}$, and a universal feature extractor f_θ shared for all domains (assuming that images from all domains are resized to the same size). In the homogeneous DG, we only need a single classifier g_ϕ that can be shared across all domains.

The proposed workflow is: (i) train a multi-domain model $g \circ f$ with \mathcal{D} , (ii) take the shared feature extractor part f_θ and use it as a fixed feature extractor for the target domain, (iii) extract features for target domain’s train set D_{N+1} and train a SVM or KNN classifier, (iv) evaluate on the testing split of the target domain, denoted as $\tilde{D}_{N+1} = \{\tilde{X}^{(N+1)}, \tilde{Y}^{(N+1)}\}$. In the case of homogeneous DG, we can also take the shared classifier f_ϕ and use the full model $g_\phi \circ f_\theta$ directly for the target domain. The goal is to perform the training in step (i) above so that the feature extractor f_θ is robust enough to perform well on any target domain without fine-tuning.

3.1. A Simple Baseline

A naive deep learning approach called aggregation (AGG) trains a single extractor to minimise the total cross-entropy

Algorithm 1 Simulating Domain Shift in Training

Input: $\mathcal{D} = \{D_1, D_2, \dots, D_N\}$

Output: θ

begin

while not converge or reach max steps **do**

Randomly split \mathcal{D} :

$\mathcal{D}_{\text{trn}} \cap \mathcal{D}_{\text{val}} = \emptyset$

$\mathcal{D}_{\text{trn}} \cup \mathcal{D}_{\text{val}} = \mathcal{D}$

for $t \in [1, 2, \dots, T]$ **do**

 Sample mini-batch D_{tr} from each $D_j \in \mathcal{D}_{\text{trn}}$

 Optimise feature extractor g_θ on D_{tr} using supervised and auxiliary loss h_ω .

 Sample mini-batch D_{val} from each $D_k \in \mathcal{D}_{\text{val}}$

 Optimise auxiliary loss h_ω on D_{val}

(CE) loss of all domains.

$$\min_{\theta, \phi_1, \phi_2, \dots, \phi_N} \sum_{D_j \in \mathcal{D}} \sum_{d_j \in D_j} \ell^{(\text{CE})}(g_{\phi_j}(f_\theta(x^{(j)})), y^{(j)}) \quad (1)$$

Here D_j is the j th domain and d_j is a mini-batch of it.

Then we fix f_θ and extract features for the training split of target domain D_{N+1} , i.e., $f_\theta(X^{(N+1)})$. With those extracted features, we can train a classifier using the pairs $\{f_\theta(X^{(N+1)}), Y^{(N+1)}\}$. Finally we test the model on the testing split of target domain $\tilde{D}_{N+1} = \{\tilde{X}^{(N+1)}, \tilde{Y}^{(N+1)}\}$.

This simple baseline surpasses many prior purpose designed DG methods as discussed in Li et al. (2017a). The key question is how to improve this naive approach, such that the trained feature extractor f_θ produces more robust features that generalise better to unseen target domains.

3.2. Simulating Domain Shift in Training

Our high-level strategy simulates domain-shift during training as illustrated in Algo. 1 and Figure 1. We use the learned feature-critic loss to guide learning on the meta-training set \mathcal{D}_{trn} , and optimise the feature-critic itself on the meta-validation set \mathcal{D}_{val} . The key idea is that training f_θ with h_ω on \mathcal{D}_{trn} should improve its performance on \mathcal{D}_{val} .

3.3. Meta-Learning an Auxiliary loss

We aim to make the model training process behave well, i.e., after each update with mini-batches from \mathcal{D}_{trn} , performance should improve for mini-batches from \mathcal{D}_{val} . This would happen to some extent without any extra effort, but we aim to enforce it by introducing a learned auxiliary loss function, denoted $\ell^{(\text{Aux})} = h_\omega$. The only requirements for h_ω are (i) it outputs a non-negative scalar (since it is a loss) and (ii) its input depends on the feature extractor’s parameter θ .

For now, we assume a suitable function h_ω (i.e., $\ell^{(\text{Aux})}$)

exists and discuss how to use it. We discuss design choices for h_ω in Sec. 3.5. With the auxiliary loss, the objective function in Eq. 1 becomes,

$$\min_{\theta, \phi_{j^s}} \sum_{D_j \in \mathcal{D}_{\text{tm}}} \sum_{d_j \in D_j} \ell^{(\text{CE})}(g_{\phi_j}(f_\theta(x^{(j)})), y^{(j)}) + \ell^{(\text{Aux})} \quad (2)$$

Taking the gradient of Eq. 2 w.r.t. θ gives two terms: (i) cross-entropy loss and (ii) auxiliary loss (recall that $\ell^{(\text{Aux})}$'s input must depend on θ , which means $\frac{\partial \ell^{(\text{Aux})}}{\partial \theta}$ generates non-zero values).

Consider two alternative updates to θ , with and without the help of $\ell^{(\text{Aux})}$. We have $\theta^{(\text{OLD})} = \theta - \alpha \frac{\partial \ell^{(\text{CE})}}{\partial \theta}$ and $\theta^{(\text{NEW})} = \theta - \alpha \frac{\partial \ell^{(\text{CE})}}{\partial \theta} - \alpha \frac{\partial \ell^{(\text{Aux})}}{\partial \theta}$. Here α is the step size of θ . If the auxiliary loss $\ell^{(\text{Aux})}$ indeed does a good job of promoting domain invariance of θ , then the latter update $\theta^{(\text{NEW})}$ exploiting $\ell^{(\text{Aux})}$ should produce a more effective feature on mini-batches from the validation domain \mathcal{D}_{val} .

Thus we train the auxiliary loss (feature-critic network) to promote this. Specifically, we optimise the parameter ω of feature-critic network as follows:

$$\max_{\omega} \sum_{D_j \in \mathcal{D}_{\text{val}}} \sum_{d_j \in D_j} \tanh(\gamma(\theta^{(\text{NEW})}, \phi_j, x^{(j)}, y^{(j)})) - \gamma(\theta^{(\text{OLD})}, \phi_j, x^{(j)}, y^{(j)}) \quad (3)$$

Here γ is a function that measures the validation domain performance (larger is better), and we discuss how to design it in Sec. 3.4. \tanh is a utility function, which converts the reward (performance gain) to utility. It reflects the commonly accepted idea concept *diminishing marginal utility*, and links $\theta^{(\text{NEW})}$ with $\theta^{(\text{OLD})}$. If \tanh and the $\theta^{(\text{OLD})}$ term are excluded in Eq. 3, it would simply maximise the validation set performance with $\theta^{(\text{NEW})}$. The reason Eq. 3 is better is that $\gamma(\theta^{(\text{OLD})}, \phi_j, x^{(j)}, y^{(j)})$ serves as a baseline, making the value range – and thus the gradient – more stable. One can understand the role of \tanh here as a smoother version of min/max-margin or a softer version of gradient clipping.

In summary, optimising the feature-critic h_ω as Eq. 3 produces a loss $\ell^{(\text{Aux})}$ that encourages the base network to extract domain agnostic features when applied in Eq. 2.

3.4. Measuring Validation Performance: Designing γ

To measure validation performance, γ can take up to four variables as input: feature extractor parameter θ , classifier parameter ϕ , data x , and label y , i.e., $\gamma(\theta, \phi, x, y)$. One simple choice is the negative classification loss, i.e.,

$$\gamma(\theta, \phi, x, y) = -\ell^{(\text{CE})}(g_\phi(f_\theta(x)), y). \quad (4)$$

Inserting Eq. 4 into Eq. 3, we have

$$\min_{\omega} \sum_{D_j \in \mathcal{D}_{\text{val}}} \sum_{d_j \in D_j} \tanh(\ell^{(\text{CE})}(g_{\phi_j}(f_\theta^{(\text{NEW})}(x^{(j)})), y^{(j)})) - \ell^{(\text{CE})}(g_{\phi_j}(f_\theta^{(\text{OLD})}(x^{(j)})), y^{(j)}) \quad (5)$$

We can now introduce the meta-loss $\ell^{(\text{Meta})}$ on \mathcal{D}_{val} to abbreviate Eq. 5 as $\min_{\omega} \ell^{(\text{Meta})}$. Note that the design of γ should reflect the demands of the testing stage. Here we choose to use classification loss because we assume the model will be deployed for a classification task eventually. An alternative choice could be a metric-based loss if we knew the final task was about retrieval. We emphasise that it is not necessary for the objective function used for the training sets (e.g., $\ell^{(\text{CE})}$) to match with the γ function.

3.5. Designing Feature-Critic h_ω

Finally, we design our feature-critic network h_ω (i.e., $\ell^{(\text{Aux})}$). Recall the requirements for such an auxiliary loss: (i) It outputs a non-negative scalar; (ii) Its input depends on θ . We note that MetaReg (Balaji et al., 2018), has a regularisation function that plays the similar role to h_ω . MetaReg proposed the following form: $h_\omega(\theta) = \sum_i \omega_i |\theta_i|$. However, this introduces the same number of parameters as θ . Doubling the number of model parameters in large modern CNNs is an expensive proposition that increases optimisation difficulty and overfitting risk.

Therefore rather than designing h_ω to take θ directly, we propose a more efficient and effective way to enable h_ω to promote the base network's generalisation. Specifically, the auxiliary loss operates on the extracted features $f_\theta(x)$. Since our auxiliary generalisation-promoting loss operates on the feature representation produced by the base network, we denote it Feature-Critic.

Denote $F = f_\theta(X^{(j)})$ as the $M \times H$ sized matrix stacking the H -dimensional features from M examples in a mini-batch from the j th domain in the virtual training set \mathcal{D}_{tm} . $h_\omega := h_\omega(F) = h_\omega(f_\theta(X^{(j)}))$. A key requirement of h_ω is that it should be permutation invariant to the rows of F , i.e., it should not make a difference if we feed images indexed $[1, 2, 3]$ or $[3, 2, 1]$. Two available choices are:

(i) The set embedding (Zaheer et al., 2017), i.e.,

$$h_\omega(F) = \frac{1}{M} \sum_{i=1}^M \text{MLP}_\omega(F_i) \quad (6)$$

where F_i denotes a row of F , and MLP is a multi-layer perceptron.

(ii) The flattened covariance matrix, i.e.,

$$h_\omega(F) = \text{MLP}_\omega(\text{Flatten}(F^T F)) \quad (7)$$

Algorithm 2 Full Algorithm

Input: $\mathcal{D} = \{D_1, D_2, \dots, D_N\}$, α , and η
Output: θ
begin

```

while not converge or reach max steps do
    Randomly split  $\mathcal{D}$ :
     $\mathcal{D}_{\text{trn}} \cap \mathcal{D}_{\text{val}} = \emptyset$ 
     $\mathcal{D}_{\text{trn}} \cup \mathcal{D}_{\text{val}} = \mathcal{D}$ 
    for  $t \in [1, 2, \dots, T]$  do
        Meta-train:
        Sample mini-batch  $d_{\text{trn}}$  from each  $D_j \in \mathcal{D}_{\text{trn}}$ 
         $\ell^{(\text{CE})}(d_{\text{trn}}) \leftarrow \text{Eq. (1)}$  //Supervised loss
         $\ell^{(\text{Aux})}(d_{\text{trn}}) \leftarrow \text{Eqs. (6 or 7)}$  //Auxiliary loss
         $\theta^{(\text{OLD})} = \theta - \alpha \nabla_{\theta} \ell^{(\text{CE})}$ 
         $\theta^{(\text{NEW})} = \theta^{(\text{OLD})} - \alpha \nabla_{\theta} \ell^{(\text{Aux})}$ 
        Meta-test:
        Sample mini-batch  $d_{\text{val}}$  from each  $D_k \in \mathcal{D}_{\text{val}}$ 
         $\ell^{(\text{Meta})}(d_{\text{val}}, \theta^{(\text{OLD})}, \theta^{(\text{NEW})}) \leftarrow \text{Eq. (5)}$ 
        //Meta-loss
        Meta-optimization:
         $\theta \leftarrow \theta - \eta (\nabla_{\theta} \ell^{(\text{CE})} + \nabla_{\theta} \ell^{(\text{Aux})})$  //Update feat.
         $\phi \leftarrow \phi - \eta \nabla_{\phi} \ell^{(\text{CE})}$  //Update classifier
         $\omega \leftarrow \omega - \eta \nabla_{\omega} \ell^{(\text{Meta})}$  //Update feature-critic
    
```

Finally, the MLP’s output should be a scalar and we place a softplus activation to make sure its output is non-negative.

3.6. Summary

Bringing all the components together, we have the full Algo. 2. To summarise, we randomly draw train/validation domains in each iteration and: Perform a putative feature extractor update on θ with and without the auxiliary Feature-Critic loss. Then generate a meta-loss based on whether or not the feature extractor update has improved performance on the validation set. Finally the feature extractor/classifier are updated using the supervised and auxiliary losses, and auxiliary loss itself is updated using the meta-loss.

4. Experiments

We evaluate our approach, first on the heterogeneous DG problem using the VD benchmark (Section 4.1), and then on the conventional homogeneous DG using Rotated MNIST and PACS (Section 4.2). Our demo code can be viewed on https://github.com/liyiyiyong/Feature_Critic.

4.1. Heterogeneous DG experiments with VD

Dataset The Visual Decathlon dataset, initially proposed for multi-domain learning (Rebuffi et al., 2017), also pro-

vides a large scale and rigorous benchmark for DG. VD contains ten diverse domains including handwritten characters, pedestrians, traffic signs, etc. The images have been pre-processed to 72×72 . To use this benchmark for DG, we aim to train a network on a subset of source domains, and produce a robust feature extractor that provides a good representation for classification in a disjoint subset of target domains. It should do so ‘out-of-the-box’, without further fine tuning. Specifically, we take the six larger datasets (CIFAR-100, Daimler Ped, GTSRB, Omniglot, SVHN and ImageNet) as source domains and hold out the four smaller datasets (Aircraft, D. Textures, VGG-Flowers and UCF101) as target domains. We use ImageNet pre-trained ResNet-18 (He et al., 2016) as the base network for all competitors. For computational efficiency, we freeze the first four blocks of ResNet-18 and only update the remaining blocks, as well as the average pooling layer, during DG training. For all methods, the final feature is used to train SVM or KNN for the target task. All methods are evaluated by both average multi-class classification accuracy in the target domains, as well as the VD-Score metric (Rebuffi et al., 2017) that rewards consistently high performance across all domains.

Competitors Few competitors can address heterogeneous DG. For these we consider AGG baseline (Eq 1), CrossGrad (Shankar et al., 2018), MetaReg (Balaji et al., 2018), and Reptile (Nichol et al., 2018). MetaReg is originally designed to produce a robust classifier given a fixed feature. We modify MetaReg to support the heterogeneous DG by (i) applying it on the feature extractor instead (as per our Feature-Critic), called MR; (ii) applying it on the final layer of feature extractor, called MR-FL. Meanwhile Reptile is designed for few-shot meta-learning. However after modifying it for multi-domain rather than multi-task meta-learning, we found it effective for heterogeneous DG.

Feature-Critic Settings We use the set embedding architecture for the critic network (Eq 6), as the covariance architecture requires too many parameters using high dimensional ResNet. During each iteration, we randomly choose four of the six source domains as meta-train, and the remaining two provide the meta-test (validation) domains. We train all components end-to-end using the AMSGrad (Reddi et al., 2018) (batch-size/per meta-train domain=64, batch-size/per meta-test domain=32, lr=0.0005, weight decay=0.0001) for 30k iterations where the lr decayed in 5K, 12K, 15K, 20K iterations by a factor 5, 10, 50, 100, respectively. Similar to MetaReg (Balaji et al., 2018), after the parameters are trained via meta-learning, we fine-tune the network on all source datasets for the final 10k iterations.

Results We first assume the full training split is available for each target domain. Table 1 shows that: (i) The original ImageNet feature transfers to novel tasks reasonably well, as observed by classic studies (Yosinski et al.,

Table 1. Recognition accuracy (%) and VD scores on four held out target datasets in Visual Decathlon using ResNet-18 extractor.

Target	SVM Classifier							KNN Classifier						
	Im.N. PT	CrossGrad	MR	MR-FL	Reptile	AGG	FC	Im.N. PT	CrossGrad	MR	MR-FL	Reptile	AGG	FC
Aircraft	16.62	19.92	20.91	18.18	19.62	19.56	20.94	11.46	15.93	12.03	11.46	13.27	14.03	16.01
D. Textures	41.70	36.54	32.34	35.69	37.39	36.49	38.88	39.52	31.98	27.93	39.41	32.80	32.02	34.92
VGG-Flowers	51.57	57.84	35.49	53.04	58.26	58.04	58.53	41.08	48.00	23.63	39.51	45.80	45.98	47.04
UCF101	44.93	45.80	47.34	48.10	49.85	46.98	50.82	35.25	37.95	34.43	35.25	39.06	38.04	41.87
Ave.	38.71	40.03	34.02	38.75	41.28	40.27	42.29	31.83	33.47	24.51	31.41	32.73	32.52	34.96
VD-Score	308	280	269	296	324	290	344	215	188	144	215	201	189	236

2014). (ii) Demonstrating the benefit of simply exploiting large datasets, the AGG baseline’s feature, trained on more than 1.39 million images across the six domains, provides strong performance. However, while it has a higher average accuracy than the ImageNet feature, AGG’s VD score is lower, reflecting its inconsistent performance. Thus obtaining consistently high scores from multi-domain training is non-trivial. Naively aggregating more diverse source domains into training can both help and hinder performance (for example, depending on if aggregated domains are particularly similar or dissimilar to a given target). Nevertheless, AGG sometimes outperforms prior purpose designed DG methods CrossGrad and MetaReg, with only Reptile producing a feature that outperforms AGG in both accuracy and VD-score metrics. (iii) Overall, our Feature-Critic (FC) method generally provides the best performance across domains and across both types of classifiers evaluated.

Although the above application scenario of heterogeneous DG is one where compute, memory or human resources rule out feature fine-tuning, another motivating scenario is where the target domain data is too sparse for effective fine-tuning. Thus we next investigate the situation if less target data is available. Specifically, we repeat the evaluation assuming that [10%, 25%, 50%, 100%] of the training split is available for SVM/KNN training. Table 2 reports target domain test accuracies under these settings. We can see that Feature-Critic provides a consistent improvement over the alternatives. Finally, we also consider a genuinely few-shot setting for the target domain. In this case we consider $K = 3, 5, 8, 10$ labelled examples per class in the target domain, and perform KNN recognition on their test sets. The results in Table 3 show that for simple similarity-based matching in a novel target domain, Feature-Critic also provides the best off-the-shelf feature representation.

In summary the Feature-Critic meta-training strategy produces a feature extractor that is generally useful for diverse target problems in an off-the-shelf feature + shallow classifier configuration. The results outperform both the standard ImageNet feature and the obvious Data Aggregation extension across a range of operating points in the target domain from the few to many-shot regime. This suggests that Feature-Critic trained feature extractors are of wide potential value in diverse applications.

4.2. Homogeneous DG experiments

4.2.1. ROTATED MNIST

Dataset and Settings Rotated MNIST (Ghifary et al., 2015) contains six domains with each corresponding to a degree of roll rotation in the classic MNIST dataset. The basic view (M0) is formed by randomly choosing 100 images each of ten classes from the original MNIST and we create 5 rotating domains from M0 with 15° rotation each in clockwise direction, denoted by M15, M30, M45, M60, and M75. Following the setting in (Shankar et al., 2018), we perform leave-one-domain-out experiments by picking one domain to hold out as the target. We compare AGG baseline, as well as CrossGrad and MetaReg. For a recognition network all competitors use the standard MNIST CNN with two conv and one FC layer as the feature network and another FC layer as the classifier. We note prior studies (Ghifary et al., 2015; Shankar et al., 2018; Deshmukh et al., 2017) did not release specific selection of digits from within MNIST, so our results do not match the numbers in those papers exactly. However, we repeat all experiments 10 times and report the mean and standard deviation of recognition accuracy.

For Feature-Critic, we train using the AMSGrad optimizer ($\text{lr}=0.001$, $\text{weight decay}=0.00005$) for 5,000 iterations. For each iteration, one meta-train and one meta-test domain are chosen randomly from the five source domains. We also use this opportunity to compare the two variants of our loss function: Feature-Critic-MLP and Feature-Critic-Flatten.

Results It can be seen from Table 4 that AGG is again a strong baseline to beat. Over ten trials of 1000 digit samples, CrossGrad and MetaReg failed to match AGG, with only Reptile matching AGG’s performance. Meanwhile, Feature-Critic performs well with both variants of the auxiliary loss network, with the set embedding (Eq. 6) performing slightly better than the covariance matrix embedding (Eq. 7).

To qualitatively visualise the results we perform PCA projections of the features in the target domain. Figure 2 shows these projections, taking as an example the M15 domain as held out. Each dot denotes an image and the colour denotes its label. We can see that Feature-Critic (take MLP style as an example) feature extractor provides improved separability in the target domain compared to the AGG baseline.

Table 2. VD recognition accuracy differences (%) against AGG with different proportions of training data available.

Data	SVM Classifier							KNN Classifier						
	AGG	Im.N. PT	CrossGrad	MR	MR-FL	Reptile	FC	AGG	Im.N. PT	CrossGrad	MR	MR-FL	Reptile	FC
10%	18.27	+1.74	+0.58	-2.38	-1.50	+1.83	+1.31	13.93	+1.15	-0.29	-3.08	+0.63	+1.01	+1.39
25%	30.10	-5.14	-2.48	-8.44	-2.48	-2.13	+0.87	23.80	-3.62	-0.06	-9.24	-3.78	-1.32	+1.48
50%	34.63	-3.55	-0.89	-7.04	-0.07	+0.37	+2.12	30.19	-5.02	-0.41	-8.87	-4.17	-0.15	+2.60
100%	40.27	-1.56	-0.24	-6.25	-1.52	+1.01	+2.02	32.52	-0.69	+0.95	-8.01	-1.11	+0.21	+2.44
Ave.	30.82	-2.13	-0.76	-6.03	-1.39	+0.27	+1.58	25.11	-2.05	+0.05	-7.30	-2.11	-0.06	+1.98

Table 3. Recognition accuracy (%) averaged over 10 test runs on VD K-shot learning.

K	Im.N. PT	CrossGrad	MR	MR-FL	Reptile	AGG	Feature-Critic
3	13.88 ± 1.82	14.01 ± 1.98	8.86 ± 1.52	13.26 ± 1.62	14.59 ± 2.30	13.75 ± 1.77	15.18 ± 2.26
5	17.63 ± 1.55	17.74 ± 1.40	12.01 ± 1.51	17.22 ± 0.84	19.17 ± 0.94	18.00 ± 1.58	19.02 ± 1.57
8	21.58 ± 1.12	20.98 ± 1.07	14.37 ± 1.35	21.61 ± 0.92	21.24 ± 1.40	21.36 ± 1.05	22.39 ± 1.00
10	23.40 ± 0.99	22.84 ± 0.89	15.61 ± 0.88	22.80 ± 0.61	23.42 ± 1.12	22.56 ± 0.84	24.23 ± 1.00
Ave.	19.12	18.89	12.71	18.72	19.61	18.92	20.21




Figure 2. Rotated MNIST. PCA projections of target domain M15 features. Left: AGG. Right: Feature-Critic. Color: Digit.

Further Analysis Figure 3 reports the loss curves of cross entropy loss, auxiliary loss, and meta-loss for Feature-Critic during training. The cross entropy loss converges to zero, as the network usually can fit the training data perfectly. The auxiliary loss fluctuates up and down for the early stage of training, and finally stabilises to a small value. Because the auxiliary loss function h_ω itself is learned, its behaviour changes with its own learning process, which explains the fluctuations, esp. for the early stage. It is more interesting to see the pattern of meta-loss, which is the performance difference of feature extractor parameterised by $\theta^{(OLD)}$ and that by $\theta^{(NEW)}$. If we select zero as a threshold, meta-loss has a clear pattern: “above zero” → “below zero” → “being zero”. This pattern is expected because: (i) For the early stage, the auxiliary loss’ parameters are randomly initialised, so it knows little about how to help generalise, thus the gradients produced by it are rather random and less likely to help. Thus $\theta^{(OLD)}$ -based model outperforms $\theta^{(NEW)}$ -based model. (ii) With the updating of ω , h_ω improves and begins to make $\theta^{(NEW)}$ better than $\theta^{(OLD)}$. During this period, gradients produced by the auxiliary loss help the model learn to generalise. (iii) For the late stage, meta-loss goes towards zero, which indicates that h_ω no longer helps (but it does not




Figure 3. Loss curves of Feature-Critic on Rotated MNIST. Left to right: CE loss, auxiliary loss and meta-loss during DG training. Top: Feature-Critic-MLP. Bottom: Feature-Critic-Flatten.

hurt either), as all of its knowledge has now been distilled into the feature extractor. The pattern of the three losses also demonstrates that, empirically, the whole algorithm converges, including the learned auxiliary loss.

4.2.2. EVALUATION ON PACS DATASET

Dataset and Settings PACS (Li et al., 2017a) is a recent object recognition benchmark for domain generalisation. PACS contains 9991 images of size 224×224 from four different domains - Photo, Art painting, Cartoon and Sketch. It has 7 categories across these domains: dog, elephant, giraffe, guitar, house, horse and person. We follow the standard protocol and perform leave-one-domain-out evaluation. Beyond this there have been two splits of PACS used in the literature. PACS was defined with a train/validation/test split within each domain. In Li et al. (2018a) models are trained on the train split alone with the validation split used for early stopping. In Balaji et al. (2018) the combined train+validation splits were used to train the models, resulting in slightly higher performance due to more data. For direct comparison with previously published results we

Table 4. Recognition accuracy (%) averaged over 10 train+test runs on Rotated MNIST.

Target	CrossGrad	MetaReg	Reptile	AGG	Feature-Critic-MLP	Feature-Critic-Flatten
M0	86.03 ± 0.69	85.70 ± 0.31	87.78 ± 0.30	86.42 ± 0.24	89.23 ± 0.25	87.04 ± 0.31
M15	98.92 ± 0.53	98.87 ± 0.41	99.44 ± 0.22	98.61 ± 0.27	99.68 ± 0.24	99.53 ± 0.27
M30	98.60 ± 0.51	98.32 ± 0.44	98.42 ± 0.24	99.19 ± 0.19	99.20 ± 0.20	99.41 ± 0.18
M45	98.39 ± 0.29	98.58 ± 0.28	98.80 ± 0.20	98.22 ± 0.24	99.24 ± 0.18	99.52 ± 0.24
M60	98.68 ± 0.28	98.93 ± 0.32	99.03 ± 0.28	99.48 ± 0.19	99.53 ± 0.23	99.23 ± 0.16
M75	88.94 ± 0.47	89.44 ± 0.37	87.42 ± 0.33	88.92 ± 0.43	91.44 ± 0.34	91.52 ± 0.26
Ave.	94.93	94.97	95.15	95.14	96.39	96.04

Table 5. Cross-domain recognition accuracy (%) on PACS using train split (Li et al., 2017b) for training.

Target	DICA	D-MTAE	DSN	TF-CNN	MLDG	DANN	CIDDG	Reptile	CrossGrad	MetaReg	AGG	FC
A	64.6	60.3	61.1	62.9	66.2	63.2	62.7	63.4	61.0	63.5	63.3	64.4
C	64.5	58.7	66.5	67.0	66.9	67.5	69.7	67.5	67.2	69.5	66.3	68.6
P	91.8	91.1	83.3	89.5	88.0	88.1	78.7	88.7	87.6	87.4	88.6	90.1
S	51.1	47.9	58.6	57.5	59.0	57.0	64.5	55.9	55.9	59.1	56.5	58.4
Ave.	68.0	64.5	67.4	69.2	70.0	69.0	68.9	68.9	67.9	69.9	68.7	70.4

Table 6. Cross-domain recognition accuracy (%) on PACS using train+validation split (Balaji et al., 2018) for training.

Target	CrossGrad	MetaReg	Reptile	AGG	Feature-Critic
A	64.84	69.82	64.35	63.77	64.89
C	67.69	70.35	70.09	66.77	71.72
P	88.48	91.07	88.78	88.62	89.94
S	57.52	59.26	59.91	57.27	61.85
Ave.	69.63	72.62	70.78	69.11	72.10

evaluate both of these settings.

The ImageNet pre-trained AlexNet (Krizhevsky et al., 2012) is used as the backbone network. Our competitors include: DICA (Muandet et al., 2013), D-MTAE (Ghifary et al., 2015), DSN (Bousmalis et al., 2016), TF-CNN (Li et al., 2017a), MLDG (Li et al., 2018a), DANN (Ganin et al., 2016), CIDDG (Li et al., 2018c), Reptile (Nichol et al., 2018), CrossGrad (Shankar et al., 2018) and MetaReg (Balaji et al., 2018). We note that DANN is designed for domain adaptation, and Reptile for few-shot learning. We re-purpose them for DG. DANN, Reptile, CrossGrad, AGG, and MetaReg in Table 5 are our implementations. The other results are taken from Li et al. (2018a), Li et al. (2018c) and Balaji et al. (2018). Among these, MetaReg makes a domain general classifier; MLDG aligns gradients to achieve a more robust optima; CrossGrad synthesises data for a new domain; DANN makes indistinguishable representations across source domains; CIDDG learns discriminative features to match the distributions across domains; Feature-Critic (FC) learns representations that generalise to new domains with a better objective since it trains a supervised loss and explicitly simulates domain shift.

Our Feature-Critic (set embedding variant) is trained with M-SGD optimizer (batch size/per meta-trian domain=32,

batch size/per meta-test domain=16, lr=0.0005, weight decay=0.00005, momentum=0.9) for 45K iterations. At each iteration, we randomly choose two of the three source domains as meta-train and the remaining one as meta-test.

Results The comparison with state-of-the-art methods on PACS dataset is shown in Table 5 and Table 6. AGG provides a hard baseline to beat as usual. Nevertheless Feature-Critic performs comparably to the best performing state of the art alternative in both settings of this benchmark.

5. Conclusion

We addressed the domain generalisation problem with a particular focus on the heterogeneous case, by meta-learning a regulariser to help train a feature extractor to be domain invariant. The resulting feature extractor outperforms alternatives for general purpose use as a fixed downstream image encoding. Evaluated on Visual Decathlon – the largest DG evaluation thus far – this suggests that Feature-Critic trained feature extractors could be of wide potential value in diverse applications. Furthermore Feature-Critic also performs favourably compared to state-of-the-art in the homogeneous DG setting. In future work we will apply Feature-Critic to other problems including RL, and explore the impact on fine-tuning target problems.

Acknowledgements

This work was produced while the first author was visiting the University of Edinburgh. This work was supported by National Natural Science Foundation of China (Grant No. 61751208), China Scholarship Council, EPSRC grant EP/R026173/1, and NVIDIA Corporation GPU donation.

References

- Balaji, Y., Sankaranarayanan, S., and Chellappa, R. Metareg: towards domain generalization using meta-regularization. In *NeurIPS*, 2018.
- Bousmalis, K., Trigeorgis, G., Silberman, N., Krishnan, D., and Erhan, D. Domain separation networks. In *NIPS*, 2016.
- Deshmukh, A. A., Sharma, S., Cutler, J. W., and Scott, C. Multiclass domain generalization. In *NIPS Workshop on Limited Labeled Data*, 2017.
- Donahue, J., Jia, Y., Vinyals, O., Hoffman, J., Zhang, N., Tzeng, E., and Darrell, T. Decaf: A deep convolutional activation feature for generic visual recognition. In *ICML*, 2014.
- Duan, L., Tsang, I.-W., and Xu, D. Domain transfer multiple kernel learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 34(3):465–479, 2012.
- Finn, C., Abbeel, P., and Levine, S. Model-agnostic meta-learning for fast adaptation of deep networks. In *ICML*, 2017.
- Ganin, Y. and Lempitsky, V. Unsupervised domain adaptation by backpropagation. In *ICML*, 2015.
- Ganin, Y., Ustinova, E., Ajakan, H., Germain, P., Larochelle, H., Laviolette, F., Marchand, M., and Lempitsky, V. Domain-adversarial training of neural networks. *The Journal of Machine Learning Research*, 17(1):2096–2030, 2016.
- Ghifary, M., Kleijn, W. B., Zhang, M., and Balduzzi, D. Domain generalization for object recognition with multi-task autoencoders. In *CVPR*, 2015.
- Gygli, M., Norouzi, M., and Angelova, A. Deep value networks learn to evaluate and iteratively refine structured outputs. In *ICML*, 2017.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *CVPR*, 2016.
- Krizhevsky, A., Sutskever, I., and Hinton, G.-E. Imagenet classification with deep convolutional neural networks. In *NIPS*, 2012.
- Li, D., Yang, Y., Song, Y.-Z., and Hospedales, T. Deeper, broader and artier domain generalization. In *ICCV*, 2017a.
- Li, D., Yang, Y., Song, Y.-Z., and Hospedales, T. Learning to generalize: Meta-learning for domain generalization. In *AAAI*, 2018a.
- Li, H., Pan, S.-J., Wang, S., and Kot, A.-C. Domain generalization with adversarial feature learning. In *CVPR*, 2018b.
- Li, Y., Tian, X., Gong, M., Liu, Y., Liu, T., Zhang, K., and Tao, D. Deep domain generalization via conditional invariant adversarial networks. In *ECCV*, 2018c.
- Li, Z., Zhou, F., Chen, F., and Li, H. Meta-SGD: Learning to learn quickly for few shot learning. In *arXiv*, 2017b.
- Long, M., Cao, Y., Wang, J., and Jordan, M. Learning transferable features with deep adaptation networks. In *ICML*, 2015.
- Long, M., Zhu, H., Wang, J., and Jordan, M. I. Unsupervised domain adaptation with residual transfer networks. In *NIPS*, 2016.
- Mottian, S., Piccirilli, M., Adjeroh, D.-A., and Doretto, G. Unified deep supervised domain adaptation and generalization. In *ICCV*, 2017.
- Muandet, K., Balduzzi, D., and Schölkopf, B. Domain generalization via invariant feature representation. In *ICML*, 2013.
- Nichol, A., Achiam, J., and Schulman, J. On first-order meta-learning algorithms. In *arXiv*, 2018.
- Niu, L., Li, W., and Xu, D. Multi-view domain generalization for visual recognition. In *ICCV*, 2015a.
- Niu, L., Li, W., and Xu, D. Visual recognition by learning from web data: A weakly supervised domain generalization approach. In *CVPR*, 2015b.
- Ravi, S. and Larochelle, H. Optimization as a model for few-shot learning. In *ICLR*, 2017.
- Razavian, A., Azizpour, H., Sullivan, J., and Carlsson, S. Cnn features off-the-shelf: An astounding baseline for recognition. In *CVPRW*, 2014.
- Rebuffi, S.-A., Bilen, H., and Vedaldi, A. Learning multiple visual domains with residual adapters. In *NIPS*, 2017.
- Rebuffi, S.-A., Bilen, H., and Vedaldi, A. Efficient parametrization of multi-domain deep neural networks. In *CVPR*, 2018.
- Reddi, S. J., Kale, S., and Kumar, S. On the convergence of adam and beyond. In *ICLR*, 2018.
- Rosenfeld, A. and Tsotsos, J. Incremental learning through deep adaptation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2018.

- Saito, K., Ushiku, Y., and Harada, T. Asymmetric tri-training for unsupervised domain adaptation. In *ICML*, 2017.
- Schmidhuber, J., Zhao, J., and Wiering, M. Shifting inductive bias with success-story algorithm, adaptive levin search, and incremental self-improvement. *Machine Learning*, 1997.
- Shankar, S., Piratla, V., Chakrabarti, S., Chaudhuri, S., Jyothi, P., and Sarawagi, S. Generalizing across domains via cross-gradient training. In *ICLR*, 2018.
- Shu, R., Bui, H. H., Narui, H., and Ermon, S. A dirt-t approach to unsupervised domain adaptation. In *ICLR*, 2018.
- Snell, J., Swersky, K., and Zemel, R. S. Prototypical networks for few shot learning. In *NIPS*, 2017.
- Sung, F., Zhang, L., Xiang, T., Hospedales, T., and Yang, Y. Learning to learn: Meta-critic networks for sample efficient learning. In *arXiv 1706.09529*, 2017.
- Sung, F., Yang, Y., Zhang, L., Xiang, T., Torr, P., and Hospedales, T. Learning to compare: relation network for few-shot learning. In *CVPR*, 2018.
- Thrun, S. and Pratt, L. (eds.). *Learning to Learn*. Kluwer Academic Publishers, 1998.
- Volpi, R., Namkoong, H., Sener, O., Duchi, J., Murino, V., and Savarese, S. Generalizing to unseen domains via adversarial data augmentation. In *NeurIPS*, 2018.
- Xu, Z., Hasselt, H., and Silver, D. Meta-gradient reinforcement learning. In *NeurIPS*, 2018.
- Yang, Y. and Hospedales, T. A unified perspective on multi-domain and multi-task learning. In *ICLR*, 2015.
- Yosinski, J., Clune, J., Bengio, Y., and Lipson, H. How transferable are features in deep neural networks? In *NIPS*, 2014.
- Zaheer, M., Kottur, S., Ravanbakhsh, S., Póczos, B., Salakhutdinov, R. R., and Smola, A. J. Deep sets. In *NIPS*, 2017.